is a function in the second stratum and is implemented by the core controllers. Based on the measurements taken by routing functions implemented on routing means 101, requirements for resource allocations are passed as triggers from the edge controllers 250 to the core controllers 260. Edge controllers 250 also send the releases of allocated resources for some other edge-to-edge links. The requirements for resource allocation and resource release enable the core controllers 260 to configure network resources. The term configuring as used herein refers to the process of capacity allocation so that end-to-end capacity requirements are met. This capacity may be defined in terms of bandwidth, data rate, bit rate, byte rate, or any other suitable measurement.

According to this embodiment, using edge-core controller links 810, correction in resource allocations may be time-coordinated between the edge controllers 250 and core controllers 260. This allows the edge controllers 250 to utilize the corrections in resource allocations as soon as they are made. As an alternative, edge controllers can wait with some guard-time. However, without time-coordination this guard-time would tend to be very high, typically in the order of several milliseconds, and would result in sub-optimal use of resources, as the corrections made would not be utilized immediately.

One mechanism for time-coordination is described in applicant's United States Patent Application Serial No. 09/286,431, filed on April 6, 1999, and titled SELF-CONFIGURING DISTRIBUTED SWITCH, now US Pat. 6,570,872, the specification of which is incorporated herein by reference. In this mechanism, edge controllers 250 send timing packets to core controllers 260 that are time stamped at the core controllers 260 and returned to the edge controllers 250. This allows the edge controllers 250 to adjust their timing according to the core controller. However, when core nodes are interconnected and edge-to-edge links are realized over multiple core-core links, time coordination becomes difficult. For example, for TDM networks, it is difficult to align time frames at the input ports of any bufferless core node, except for a core node that receives data exclusively from source edge nodes. Even if misalignment problem is solved, there is still the problem of time-slot vacancy mismatch, which is difficult to solve without time-slot interchange, which requires the use of buffers. Therefore, according to this embodiment, the network may preferably limit the number of core nodes that together implement a single edge-to-edge link, the number preferably being 1 or 2. This allows the time-

MMM 11/23/2006